What is claimed is:

1. A system for speech unit selection comprising:

    a large speech database referencing speech waveforms and associated symbolic prosodic features, wherein the speech database is accessed by speech waveform designators, at least one designator being associated with a sequence of one or more diphones; and

    a speech waveform selector, in communication with the speech database, that selects based, at least in part, on the symbolic prosodic features stored in the speech database, waveforms referenced by the speech database.

2. A system according to claim 1, wherein the speech waveform selector uses criteria that favor approximately equally all waveform candidates having low level prosody features within a target range determined as a function of high level linguistic features.

3. A system for speech unit selection comprising:

    a large speech database referencing speech waveforms;

    a speech waveform selector, in communication with the speech database, that selects waveforms referenced by the speech database using criteria that, at least in part, favor (i) waveform candidates based directly on high level prosody features, and (ii) approximately equally all waveform candidates having low level prosody features within a target range determined as a function of high level linguistic features.

4. A system according to claim 2 or 3, wherein the criteria include a first requirement favoring waveform candidates having pitch within a target range determined as a function of high level linguistic features.

5. A system according to claim 2 or 3, wherein the criteria include a second requirement favoring waveform candidates having a duration within a target range determined as a function of high level linguistic features.

6. A system according to claim 2 or 3, wherein the criteria include a third requirement favoring waveform candidates having coarse pitch continuity within a target range determined as a function of high-level linguistic features.

7.   A system according to claim 2 or 3, wherein the synthesizer operates to select among waveform candidates without recourse to specific target duration values or specific target pitch contour values over time.